arXiv:1608.05619v1 [cs.PL] 19 Aug 2016

Symbolic Abstract Contract Synthesis in a Rewriting Framework *

María Alpuente[1], Daniel Pardo[1], and Alicia Villanueva[1]

DSIC, Universitat Politècnica de València
Camino de Vera s/n, 46022 Valencia, Spain
{alpuente,daparpon,villanue}@dsic.upv.es

Abstract We propose an automated technique for inferring software contracts from programs that are written in a non-trivial fragment of C, called KernelC, that supports pointer-based structures and heap manipulation. Starting from the semantic definition of KernelC in the $\mathbb{K}$ framework, we enrich the symbolic execution facilities recently provided by $\mathbb{K}$ with novel capabilities for assertion synthesis that are based on abstract subsumption. Roughly speaking, we define an abstract symbolic technique that explains the execution of a (modifier) C function by using other (observer) routines in the same program. We implemented our technique in the automated tool KindSpec 2.0, which generates logical axioms that express pre- and post-condition assertions by defining the precise input/output behavior of the C routines.

Keywords: contracts, automatic inference, symbolic execution, formal semantics, abstract subsumption.

1 Introduction

Checking software contracts [17] is one of the most promising techniques for achieving software reliability. Contracts essentially consist of requirements that are imposed on the arguments and result values when functions are invoked. Given its interest, considerable effort has recently been invested towards giving automatic support for equipping programs with extensive contracts, yet the current contract inference tools are still often unsatisfactory in practice [8].

This paper describes a symbolic inference system that synthesizes contracts for heap-manipulating programs that are written in a non-trivial fragment of C, called KernelC [10], which includes functions, I/O primitives, dynamically allocated structures, and pointer manipulation. By automating the tedious and time-consuming process of generating contracts, programmers can reap the benefits of assertion–based debugging and verification methods with reasonable effort.

Given a program P, the contract discovery problem is generally described as the problem of inferring a likely specification for every function m in P that uses I/O primitives and/or modifies the state. The specifications that we aim to infer consist of logical assertions that characterize the function behavior and that are expressed as method

* This work has been partially supported by the EU (FEDER) and Spanish MINECO under grants TIN2015-69175-C4-1-R and TIN2013-45732-C4-1-P, and by Generalitat Valenciana PROMETEOII/2015/013. Daniel Pardo was supported by FPU-ME grant FPU14/01830.

pre-conditions (imposed on the arguments) and post-conditions (relating the arguments and the result for a method).

In [1], a preliminary specification inference technique was proposed that is based on the classification scheme for data abstractions developed in [16], where a function (method) may be either a *constructor*, a *modifier*, or an *observer*. The intended behavioral specification of any *modifier* function m of P is expressed as a set of logical assertions that characterize the pre- and post-states of the m execution by using the *observer* functions in P. For instance, for the case of a modifier function `push` that adds the element `x` into a given bounded stack `q` (and assuming the traditional meaning for the observer functions `top`, `isfull`, and `size`), a typically expected axiom could be `isfull(q)=0` $\wedge$ `size(q)=n` $\Rightarrow$ `top(q)=x` $\wedge$ `size(q)=n+1`.

The inference technique of [1] relies on symbolic execution (SE) [15], a well-known program analysis technique that allows programs to be executed using *symbolic* input values instead of actual (concrete) data so that the program execution manipulates symbolic expressions involving the symbolic values. More precisely, for each pair (s,s') of initial and final states in the symbolic execution of m, an implicative axiom $p \Rightarrow q$ is synthesized where both the antecedent p and the consequent q are expressed in terms of the (sub-)set of program observers that *explain* s and s'. This is achieved by analyzing the results of symbolically executing each observer method o from initial configurations that contain a symbolic characterization of s and s'. The symbolic infrastructure employed in [1] was built on top of the (rewriting logic) semantic framework $\mathbb{K}$, which greatly facilitates the development of executable semantics of programming languages and related formal analysis techniques [19]. Unfortunately, it was developed by reusing spare features of a formal verifier for KernelC (called MatchC) that was formerly provided within $\mathbb{K}$ but is currently unsupported. On the other hand, the underlying methodology in [1] was rather limited since a fixed threshold for loop unrolling was imposed in order to avoid non-termination risks. In [2], we switched to a recent, native extension of $\mathbb{K}$ that supports symbolic execution by language transformation [4]. However, the methodology in [2] inherited the loop unrolling strategy based on depth bounds from [1].

In this work, we improve the inference power of [1,2] by endowing $\mathbb{K}$'s symbolic execution with modern subsumption techniques based on approximation [3] and lazy initialization [14]. The fact that this symbolic infrastructure is much more flexible and (potentially) language-independent allows us to define a generic, more accurate, easily maintainable and robust framework for the inference of program contracts that could be adapted to other languages defined within the $\mathbb{K}$ framework with negligible effort. We summarize our contributions as follows.

1. A symbolic algorithm that synthesizes contracts for heap-manipulating code while coping with infinite computations. This is done by
   (a) augmenting $\mathbb{K}$'s symbolic execution with lazy initialization and a widening operator based on abstract subsumption (in Section 4), and
   (b) synthesizing method pre- and post-conditions by means of a contract inference algorithm that explains the (initial and final) abstract symbolic execution states by using the program observers (in Section 5).

   Because of the abstraction, some inferred axioms for method m cannot be guaranteed to be correct and are kept apart as "candidate" (or overly general) axioms. A contract refinement algorithm is then formalized that tries to falsify them by checking whether

an input call to m that satisfies the axiom antecedent ends in a final state that does not satisfy the given consequent.
2. The proposed inference technique is implemented in the KindSpec 2.0 system, which builds on the capabilities of the SMT solver Z3 [18] to simplify the axioms. Also, the inferred contracts are given a compact representation that abstracts the user from any implementation details.

2 Method Specification: A Running Example

By abuse, we use the standard terminology for contracts of object-oriented programming and speak of *methods* when we refer to KernelC functions. Like many state-of-the-art formal specification approaches, we assume to be working in a contract-based setting [17], where the granularity of specification units is at the level of one method. Our inference technique relies on the classification scheme developed for data abstractions in [16], where a function (method) may be either a *constructor*, a *modifier*, or an *observer*. A constructor returns a new data object from scratch (i.e., without taking the object as an input parameter). A modifier alters an existing object (i.e., it changes the state of one or more of the data attributes in the instance). An observer inspects the object and returns a value characterizing one or more of its state attributes. Since the C language does not enforce data encapsulation, we cannot assume purity of any function; thus, we do not assume the traditional premise that states that observer functions do not cause side effects. In other words, any function can potentially be a modifier and we simply define an observer as any function whose return type is different from `void`.

Let us introduce the leading example that we use to describe our inference methodology: a KernelC implementation of an abstract datatype for representing sets by using linked lists. The example is composed of 7 methods: one constructor (`new`), one modifier (`insert`), and five observers (`isnull`, `isempty`, `isfull`, `contains`, and `length`). Note that the observers in this program do not modify any program objects, even if purity of observers is not required in our framework. As is usual in C, logical observers return value `1` (resp. `0`) to represent the traditional boolean value *true* (resp. *false*).

Example 1. Consider the program fragment given in Figure 1 (the full program code can be found in Appendix A), where we define set operations over a data structure (`struct set`) that records the number of elements contained in the set (field `size`), the maximum number of elements that can be held (field `capacity`), and a pointer to a list that stores the set elements (field `elems`). Each node of the list is a record data structure (`struct lnode`) that contains an integer value (field `value`) and a pointer to the subsequent list element (field `next`).

A call `insert(s,x)` to the `insert` function proceeds as follows: it first checks that the pointer `s` to the set structure is different from `NULL`, that the set is not full, and that `x` is not in the set yet. Then, a new list node `*new_node` is allocated, filled with the value `x`, and inserted as the first element of the list. Also, the size of the set is increased by 1 and the call returns `1`; otherwise, `0` is returned and `s` is not modified.

The following observers return `0` unless explicitly stated otherwise. `isnull(s)` returns `1` if the pointer `s` references to `NULL` memory; `isempty` returns `1` if `s` is initialized but `elems` is `NULL`; `isfull(s)` returns `1` if the size of `s` is greater than or equal to its capacity; and `contains(s,x)` returns `1` if the value `x` is found in `s`. The function `length(s)` incrementally counts the number of elements (nodes) in the set `s` by traversing the list `s->elems` and returns this number, or it returns `0` if the set `s` pointer is `NULL`.

```
1  #include <stdlib.h>
2
3  struct lnode{
4   int value;
5   struct lnode *next;
6  };
7  struct set {
8   int capacity;
9   int size;
10  struct lnode *elems;
11 };
12 struct set* new(int capacity) {...}
13
14 int insert(struct set *s, int x) {
15  struct lnode *new_node;
16  struct lnode *end_node;
17  struct lnode *n;
18
19  if(s==NULL)
20   return 0; /* NULL set */
21  if(s->size >= s->capacity)
22   return 0; /* no space left */
23
24  end_node = s->elems;
25  n = end_node;
26  while(n != NULL) {
27   if(n->value == x)
28    return 0; /* x already in the set */
29   end_node = n;
30   n = n->next;
31  }
32
33  new_node = (struct lnode*) malloc(sizeof(struct
        lnode));
34  if(new_node == NULL)
35   return 0; /* no memory left */
36  new_node->value = x;
37  new_node->next = s->elems;
38
39  s->elems = new_node;
40  s->size += 1;
41  return 1; /* element added */
42 }
43
44 int isnull(struct set *s) {...}
45 int isempty(struct set *s) {...}
46 int isfull(struct set *s) {...}
47 int contains(struct set *s, int x) {...}
48 int length(struct set *s) {...}
```

Figure 1. Fragment of the KernelC implementation of a set datatype.

From the source code of the program, for each modifier function m, we aim to synthesize, a contract of the form $< P, Q, \mathcal{L} >$ where P is the method precondition, Q is the method postcondition, and $\mathcal{L}$ is the set of program locations (local variables, data-structure pointers and fields, and method parameters) that are (potentially) affected by the method execution. We first compute a set of implication formulas $p \Rightarrow q$, where p and q are conjunctions of equations $l = r$. The left-hand side l of each equation can be either 1) a call to an observer function, and then r represents the return value of that call; or 2) the keyword `ret`, and then r represents the value returned by the modifier function m being observed. Then, given the set of implication formulas $\{p_1 \Rightarrow q_1, \ldots, p_n \Rightarrow q_n\}$, P is defined as $p_1 \vee \ldots \vee p_n$, the postcondition Q is the formula[1] $(p_1 \Rightarrow q_1) \wedge \ldots \wedge (p_n \Rightarrow q_n)$, and the elements of $\mathcal{L}$ refer to locations whose value might be affected by the execution of m, that is, all memory locations of the pre-state that do not belong to the set $\mathcal{L}$ remain allocated and are left unchanged in the post-state. The set $\mathcal{L}$ itself is interpreted in the pre-state and is necessary for sound usage of contracts.

Example 2. The intended postcondition Q for the modifier function `insert(s,x)` of Example 1 contains five axioms (each one given by an implication), which are shown in Figure 2. We adopt the standard primed notation to distinguish variable values after the execution of the method from their value before the execution.

The first axiom can be read as: if the outcome of `isnull(s)` is 1 before the call to `insert(s,x)`, then, after execution, the set is still null and the value returned by `insert(s,x)` is 0, which means that the element `x` was not inserted.

The last axiom can be read as: if the set is neither null, full nor empty and there is no node in the list with value `x`, then, after execution, the set remains non-null and non-empty, the value `x` is now in the set, the length is increased by 1, and the call to `insert(s,x)` returns 1, which denotes a successful insertion.

[1] This is similar to the idea of contracts with *named behaviors* as provided in the ACSL contract specification language for C [6].

$$
\begin{aligned}
(\,\texttt{isnull(s)} = 1\,) &\Rightarrow (\,\texttt{isnull(s')} = 1 \ \wedge\ \texttt{ret} = 0\,)\\
(\,\texttt{isfull(s)} = 1\,) &\Rightarrow \begin{pmatrix}\texttt{isfull(s')} = 1 \ \wedge\ \texttt{contains(s', x)} = \texttt{contains(s, x)} \ \wedge\\ \texttt{length(s')} = \texttt{length(s)} \ \wedge\ \texttt{ret} = 0\end{pmatrix}\\
(\,\texttt{contains(s, x)} = 1\,) &\Rightarrow (\,\texttt{contains(s', x)} = 1 \ \wedge\ \texttt{length(s')} = \texttt{length(s)} \ \wedge\ \texttt{ret} = 0\,)\\
(\,\texttt{isempty(s)} = 1 \ \wedge \texttt{isfull(s)} = 0\,) &\Rightarrow \begin{pmatrix}\texttt{isempty(s')} = 0 \ \wedge\ \texttt{contains(s', x)} = 1 \ \wedge\\ \texttt{length(s')} = \texttt{length(s)} + 1 \ \wedge\ \texttt{ret} = 1\end{pmatrix}\\
\begin{pmatrix}\texttt{isnull(s)} = 0 \ \wedge \texttt{isempty(s)} = 0 \ \wedge\\ \texttt{isfull(s)} = 0 \ \wedge \texttt{contains(s, x)} = 0\end{pmatrix} &\Rightarrow \begin{pmatrix}\texttt{isnull(s')} = 0 \ \wedge\ \texttt{isempty(s')} = 0 \ \wedge\\ \texttt{contains(s', x)} = 1 \ \wedge\ \texttt{length(s')} = \texttt{length(s)} + 1 \ \wedge\\ \texttt{ret} = 1\end{pmatrix}
\end{aligned}
$$

Figure 2. Expected postcondition axioms for the `insert` method

3 The (symbolic) $\mathbb{K}$ Framework

$\mathbb{K}$ is a rewriting-based framework for engineering language semantics [19]. Provided that the syntax and semantics of a programming language are formalized in the language of $\mathbb{K}$, the system automatically generates a parser, an interpreter, and formal analysis tools such as model checkers and deductive theorem provers. Complete formal program semantics are currently available in $\mathbb{K}$ for Scheme, Java 1.4, JavaScript, Python, Verilog, and C among others [19].

A language definition in $\mathbb{K}$ consists of three parts: the BNF language syntax, the structure of program configurations, and the semantic rules. Program configurations are represented in $\mathbb{K}$ as potentially nested structures of labeled *cells* (or containers) that represent the program state. Similarly to the classic operational semantics, program configuration cells include a computation stack or continuation (named k), one or more environments (env, heap), and a call stack (stack) among others, and are represented as algebraic datatypes in $\mathbb{K}$.

The part of the $\mathbb{K}$ program configuration structure for the KernelC semantics that is relevant to this work is $\langle\, \langle \mathrm{K} \rangle_{\mathsf{k}} \langle \mathrm{Map} \rangle_{\mathsf{env}} \langle \mathrm{Map} \rangle_{\mathsf{heap}} \,\rangle_{\mathsf{cfg}}$, where the env cell is a mapping of variable names to their memory positions, the heap cell binds the active memory positions to the actual values (i.e., it stores information about pointers and data structures), and the k cell represents a stack of computations waiting to be run, with the left-most element of the stack being the next computation to be undertaken. For example, the configuration

$$\langle\, \langle \mathrm{tv}(\mathit{int}, 0) \rangle_{\mathsf{k}} \langle \mathtt{x} \mapsto \mathsf{\&x} \rangle_{\mathsf{env}} \langle \mathsf{\&x} \mapsto \mathrm{tv}(\mathit{int}, 5) \rangle_{\mathsf{heap}} \,\rangle_{\mathsf{cfg}} \tag{1}$$

models the final state of a computation whose return value is the integer 0 (stored in the k cell, which contains the current code to be run), while program variable `x` (stored in the env cell) has the integer value 5 (stored in the memory address given by &x in the heap cell). The symbol tv is a language construction aimed to encapsulate typed values. Variables representing symbolic memory addresses are written in sans-serif font preceded by the & symbol.

The semantic rules in $\mathbb{K}$ state how configurations (terms) evolve throughout the computation. A useful feature of $\mathbb{K}$ is that «rules only need to mention the minimum part of the configuration that is relevant for their operation».

For symbolic reasoning, $\mathbb{K}$ uses a particular class of first-order formulas with equality (encoded as boolean non–ground terms with constraints over them). These formulas, called *patterns*, specify those configurations that match the pattern algebraic structure and that satisfy its constraints. For instance, the pattern

$$\left\langle \begin{array}{c} \langle \mathrm{tv}(int, 0) \rangle_{\mathsf{k}} \\ \langle \cdots\ \mathtt{x} \mapsto \&\mathsf{x}, \mathtt{s} \mapsto \&\mathsf{s}\ \cdots \rangle_{\mathsf{env}} \\ \langle \cdots\ \&\mathsf{s} \mapsto (\mathtt{size} \mapsto ?\mathsf{s.size}, \mathtt{capacity} \mapsto ?\mathsf{s.capacity})\ \cdots \rangle_{\mathsf{heap}} \end{array} \right\rangle_{\mathsf{cfg}}$$
$$\left\langle\ \&\mathsf{s} \neq \mathsf{NULL} \wedge ?\mathsf{s.size} \geq ?\mathsf{s.capacity}\ \right\rangle_{\mathsf{path-condition}}$$

specifies the configurations satisfying that:1) the k cell only contains the integer value 0; 2) in the env cell, program variable `x` (in typographic font) is associated to the memory address &x while `s` is bound to the pointer &s; and 3) in the heap cell, the field `size` of (the data structure pointed by) &s (resp. its `capacity` field) contains the symbolic value[2] ?s.size (resp. ?s.capacity). Additionally, &s is not null and the value of its `size` field is greater than or equal to its `capacity` field.

Since patterns allow logical variables and constraints over them, by using patterns, the $\mathbb{K}$ execution principle (which is based on term rewriting) becomes *symbolic execution*. Unlike concrete execution where the path taken is determined by the input, in symbolic execution the program can take any feasible path and each possible path is associated to a *path condition*, which represents the conditions that input values have to satisfy in order to follow that path. The path condition is formed by constraints that are gathered along the path taken by the execution to reach the current program point, so each symbolic execution path stands for many actual program runs (in fact, for exactly the set of runs whose concrete values satisfy the logical constraints).

Symbolic execution in $\mathbb{K}$ relies on an automated transformation of $\mathbb{K}$ configurations and $\mathbb{K}$ rules into corresponding symbolic $\mathbb{K}$ configurations (i.e., patterns) and symbolic $\mathbb{K}$ rules that capture all required symbolic ingredients: symbolic values for data structure fields and program variables; path conditions that constrain the variables in cells; multiple branches when a condition is reached during execution, etc [4]. The transformed, symbolic rules define how symbolic configurations are rewritten during computation. Roughly speaking, by symbolically executing a program statement, the configuration cells are updated by mapping fields and variables to new symbolic values that are represented as symbolic expressions, while the path conditions (stored in a new path-condition cell) are correspondingly updated at each branching point.

In [2], an inference procedure for KernelC programs was defined using the $\mathbb{K}$ symbolic execution infraestructure described above. In order to avoid the exponential blowup that is inherent to path enumeration, the symbolic procedure of [2] follows the standard approach of exploring loops up to a specified number of unfoldings. This ensures that symbolic execution ends for all explored paths, thus delivering a finite (partial) representation of the program behavior [10]. In the following, given a method call $m(args)$ and an initial path condition ϕ, and assuming an unspecified unrolling bound for loops, we denote by $\mathrm{SE}(m(args)\{\phi\})$ the symbolic execution of method m with input arguments $args$ as described in [2], which returns the set of leaves (patterns) of the symbolic execution tree for m under the constraints given by ϕ. For any function f, by $f(args)\{\phi\}$, we represent the $\mathbb{K}$ pattern $\langle\langle f(args)\rangle_{\mathsf{k}}\ \cdots\rangle_{\mathsf{cfg}}\langle\phi\rangle_{\mathsf{path-condition}}$ that is built by inserting the call $f(args)$ at the top of the k cell and by initializing the path condition cell with ϕ.

[2] Symbolic values are preceded by a question mark.

4 Improving Symbolic Execution in $\mathbb{K}$

In this section, we extend $\mathbb{K}$'s symbolic execution machinery with lazy initialization techniques and abstract subsumption checking in order to support the synthesis of contracts for methods that require refined loop finitization and C pointer dereference and initialization.

Lazy initialization. Structured datatypes (`struct`) in C are aggregate types that define non-empty sets of sequentially allocated *member objects*[3], called fields, each of which has a name and a type. In our symbolic setting, pointer arithmetics and memory layout of C programs are abstracted by:1) operating with *symbolic addresses* instead of concrete addresses, and 2) mapping each structure object into a single element of the heap cell that groups all object fields (and associated values). A specific field is then accessed by combining the identifiers of both the structure object and the field name.

A critical point in the symbolic execution of C programs is the *undefinedness* problem that occurs when accessing uninitialized memory addresses. We adapt the lazy initialization approach of [14] to our setting as follows: when a symbolic address (or address expression) is accessed for the first time, we force SE to initialize the memory object that is located at the given address. This means that the mapping in the heap cell is updated by assigning a new symbolic value (given by the very name of the symbolic address of the accessed field) that symbolically represents the assumptions made on the dynamic data structure. Actually, when symbolic execution accesses potentially uninitialized memory positions, two cases are considered: the case in which the memory is initialized and it stores an object of its respective type; and the case in which the memory stores a null pointer.

To keep track of the constraints that are introduced by the lazy initialization, a new cell $\langle\rangle_{\text{init-heap}}$ is added to the configuration that represents the initialization assumptions on the heap memory at a given program point. In other words, at every leaf of the symbolic execution tree, the init-heap cell records the symbolic initial heap that leads to the given final symbolic configuration.

Symbolic subsumption. Symbolic execution traditionally undergoes non-termination problems in the presence of loops or recursion: the exhaustive exploration of all program paths is unaffordable because the search space may be infinite and, consequently, the number of symbolic execution paths may be unbounded. A classical solution (used in [1,2]) is to establish a *bound* to the depth of the symbolic execution tree by specifying the maximum number of unfoldings for each loop and recursive function. However, the completeness of the symbolic analysis is highly dependent on the chosen threshold, and it is not generally possible to ascertain the optimal number of iterations that *subsume* all possible behaviors by inspecting the source code.

The abstract subsumption approach of [3] determines the length of the symbolic execution paths in a dynamic way. Intuitively, symbolic execution with abstract subsumption checking proceeds as standard symbolic execution, except that before entering a loop, it is checked that the current (abstract) state has not already been explored; otherwise, the execution of the loop stops. Supporting this check does not require whole execution paths to be recorded; only symbolic states that correspond to the evaluation of loop guards need to be recorded.

[3] An object in C is a region of data storage in the execution environment.

An algorithm for symbolic subsumption that naturally transfers to our framework is given in [3]. Let us augment symbolic program configurations C into *program states* $S = \langle C, i \rangle$ by giving the configuration pattern C a program counter i that corresponds to the line number in the source code of the subsequent instruction to be executed, or the `return` statement if the configuration C is final. Also, let us represent the conjunction of all constraints over the symbolic values of primitive-type variables and structure fields expressed[4] in the env, heap, and path-condition cells of pattern C in S, called *state constraint*, by $SC(S)$. By using the subsumption algorithm, we can decide state subsumption $S_2 \sqsubseteq S_1$ by simply checking that: 1) S_1 and S_2 have the same program counter; 2) the symbolic heap in S_2 is subsumed by the symbolic heap in S_1 (i.e., the set of all possible program heaps whose concrete shape and values match the heap constraints in S_1 includes the set of all concrete program heaps that satisfy the constraints in S_2); and 3) $SC(S_2) \Rightarrow SC(S_1)$.

Abstract subsumption. Symbolic execution with state subsumption is obviously not guaranteed to terminate. In order to ensure termination and improve scalability of symbolic execution, we enhance symbolic state subsumption checking by means of abstract interpretation [3]. We abstract both primitive domains and heaps by using the abstraction function α proposed in [3]. The idea for heap abstraction is to apply a shape transformation that collapses two or more nodes into a *summary node*. Nodes can be collapsed when they are in a sequence and can only be accessed by traversing all their predecessors (i.e., each node is only pointed to by its preceding node in the sequence).

To make the visualization of symbolic heap abstraction easier, we also adapt to our symbolic setting a classical graphical representation for heaps based on UML object diagrams, where *null* nodes are rendered as ellipses, uninitialized nodes are drawn as clouds, and references are depicted as arrows.

Example 3. Node S_{27} of Figure 3 illustrates shape abstraction for the given state. The circled nodes are abstracted into a summary node. Then, the first node of the list points to this new summary node and in turn the summary node points to the node referenced by `&end_node`. Moreover, the valuation for the field `value` of the summary node (identified by e_3) is $e_3 = ?\mathsf{v}_0 \vee e_3 = ?\mathsf{v}_1$.

Given the symbolic state S, we define the abstraction $S^\sharp = \alpha(S)$. Then, the abstract symbolic subsumption relation $S_2 \sqsubseteq^\sharp S_1$ is given by $S_2^\sharp \sqsubseteq S_1^\sharp$.

4.1 Symbolic execution with abstract subsumption

The symbolic execution with abstract subsumption (and lazy initialization) of a given method m with arguments *args* and initial path condition ϕ, written $\mathrm{SE}^\sharp(m(args)\{\phi\})$, is defined as an approximation of the SE mechanism of [2] where, each time a symbolic state S_2 is visited that corresponds to a recursive call or loop guard evaluation with the same program counter as a previously visited state S_1, the abstract subsumption $S_2 \sqsubseteq^\sharp S_1$ is checked; if the test succeeds, then the loop or recursive function stops, and the execution flow proceeds to the subsequent instruction.

[4] By abuse, we assume a logical constraint representation $x_1 = v_1 \wedge \ldots \wedge x_n = v_n$ of the symbolic heap $\{x_1 \mapsto v_1, \ldots, x_n \mapsto v_1\}$, where every x_i references a field of a heap data object, whereas for the environment each x_i refers to a primitive-type program variable.

Example 4. The uncontrolled symbolic execution SE of the function `insert(s,x)` from Example 1 generates an infinite state space. In contrast, $\text{SE}^\sharp$ terminates after three iterations of the loop. Figure 3 illustrates the fragment of the symbolic execution tree for `insert(s,x)` where the subsumption between two abstract states is exposed. The state (S_{18}) corresponds to the state where the loop guard is to be checked for the third time. This requires evaluating n, which points to an uninitialized node; hence, lazy initialization is applied. This results in two children, S_{19} and S_{21}, with the same program counter because the guard has not been evaluated yet. The left child S_{19} corresponds to the case when the loop guard is not satisfied and the loop is exited, whereas the right child S_{21} represents entering the loop iteration.

Program counter 29 is reached again at state S_{27} in the right branch after lazy initialization, and then the abstract subsumption check $S_{27} \sqsubseteq^\sharp S_{21}$ succeeds.

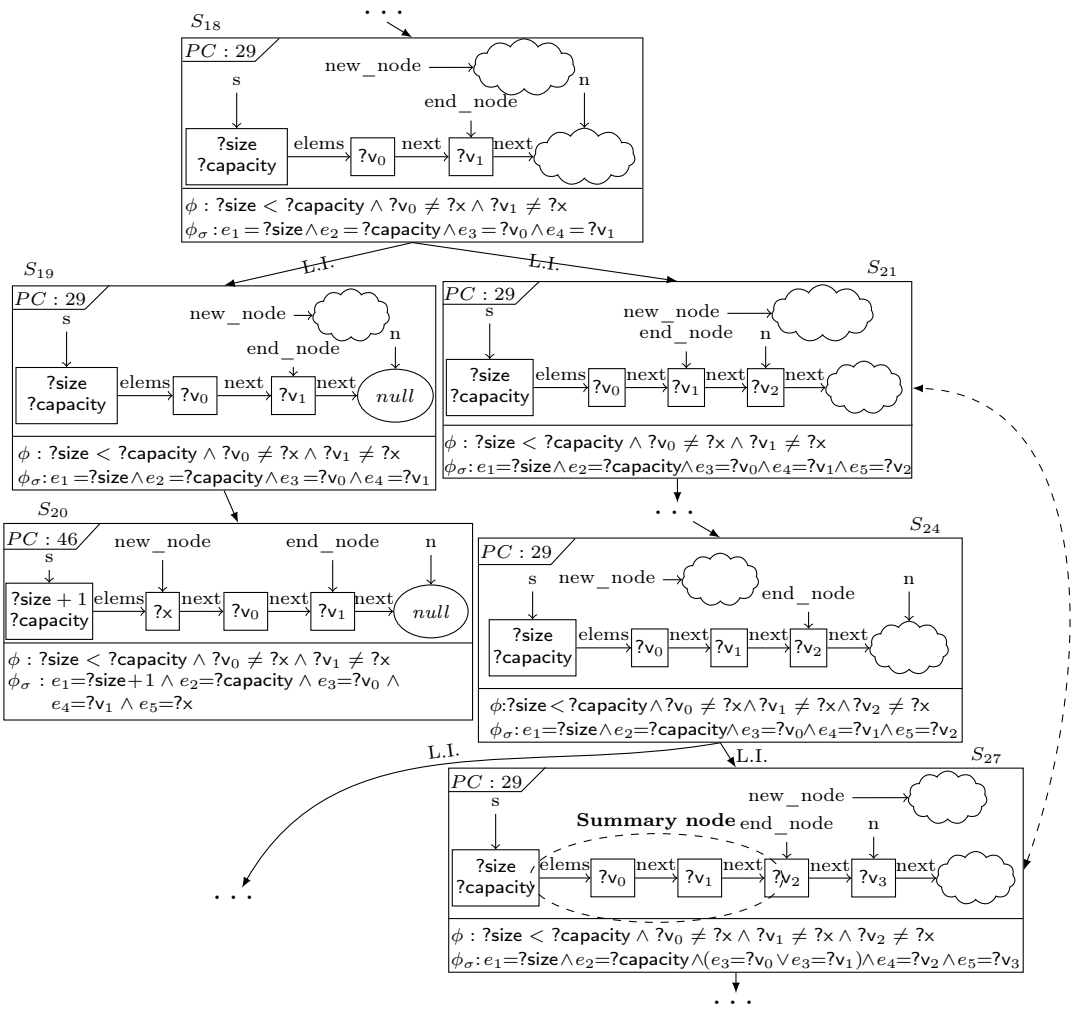

Figure 3. Fragment of the abstract symbolic execution of `insert(s,x)`

Let $\text{SE}^\sharp(f(args)\{\phi\})$ return the set of final patterns obtained from the abstract symbolic execution of the pattern $f(args)\{\phi\}$ (i.e., the leaves of the deployed abstract

symbolic execution tree). We assume appropriate abstractions are defined to ensure termination of $\text{SE}^\sharp$. A new, (abstract subsumption) cell $\langle\rangle_{\mathsf{aSubFlag}}$ identifies with a *true* value those final abstract configurations ending any branch that was folded (at some intermediate configuration) by the application of abstract subsumption. This is used for the inference process to distinguish the inferred axioms that are ensured to hold (because no approximation was done to extract them) from the *plausible*, candidate axioms that are not demonstrably correct because of the potential precision loss caused by the abstraction. Furthermore, in order to obtain the set of locations that may be affected by the execution of f (the component $\mathcal{L}$ of the contract), those locations have to be harvested during symbolic execution. To this end, we add a new cell $\langle\rangle_{\mathsf{locations}}$ to the symbolic engine of $\mathbb{K}$. Then, whenever a program location is overwritten, it is recorded in the new cell locations. At the end of the symbolic execution, the program locations recorded for each final configuration in their respective locations cells are all joined by union to obtain a global set with every program location that is potentially modifiable by a call to the function f. Thus, assignable locations are obtained as a by-product of the symbolic execution.

5 Inference Algorithm

Let us introduce the basic notions that we use in our formalization. Given an input program P, we distinguish the set of observers $\mathcal{O}$ and the set of modifiers $\mathcal{M}$ in P. A function can be considered to be an *observer* if it explicitly returns a value, whereas any method can be considered to be a *modifier*. Thus, the set $\mathcal{O} \cap \mathcal{M}$ is generally non-empty.

Algorithm 1 Specification Inference

Input: $m \in \mathcal{M}$: a modifier function with arity n
Output: *contract* : a specification contract for m

1: $root := m(\overline{a_n})$;
2: $\mathcal{F} :=\text{SE}^\sharp(root\{\texttt{true}\})$;
3: $P := false$; $Q := true$; $Q^\sharp := true$; $\mathcal{L} := \emptyset$;
4: **for all** $F \in \mathcal{F}$, with $F = \langle\langle v\rangle_{\mathsf{k}}\langle\varphi\rangle_{\mathsf{init-heap}} \cdots\rangle_{\mathsf{cfg}}\langle\phi\rangle_{\mathsf{path-condition}}\langle\sharp\rangle_{\mathsf{aSubFlag}}\langle L\rangle_{\mathsf{locations}}$ **do**
5: $p := explain(I, \overline{a_n})$, where $I = \langle\langle root\rangle_{\mathsf{k}}\langle\varphi\rangle_{\mathsf{heap}} \cdots\rangle_{\mathsf{cfg}}\langle\phi\rangle_{\mathsf{path-condition}}$;
6: $q := explain(F, \overline{a_n}) \wedge (ret = v)$;
7: $P := P \vee p$;
8: $ax := (p \Rightarrow q)$;
9: **if** $\sharp$ **then** $Q := Q \wedge ax$; **else** $Q^\sharp := Q^\sharp \wedge ax$;
10: $\mathcal{L} := \mathcal{L} \cup \{L\}$;
11: **end for**
12: **return** $< P, refine(Q, Q^\sharp), \mathcal{L} >$;

Our specification inference methodology is formalized in Algorithm 1. Let $\overline{a_n}$ denote the list of fresh symbolic variables $a_1, \ldots, a_n$. First, the *modifier* method of interest m is symbolically executed with argument list $\overline{a_n}$ and empty path constraint `true`, and the set $\mathcal{F}$ of final configurations is retrieved from the leaves of the abstract symbolic execution tree. For each final configuration the corresponding path condition ϕ is simplified by calling the automated theorem prover Z3.

After initializing the contract components (Line 3), we proceed to compute one axiom for each (abstract) symbolic configuration F in $\mathcal{F}$. First, the premise p of the axiom $p \Rightarrow q$

is computed (Line 5) by means of the function $explain(I, as)$ originally proposed in [1]. This function receives as argument the pattern I, which expresses the initial symbolic configuration leading to F in the execution tree for m (i.e., a variant of the initial configuration for $m(\overline{a_n})$ that is obtained by assuming the constraints φ and ϕ in the corresponding init-heap and path-condition cells of I). Roughly speaking, by means of a conjunction of equations $explain(I, \overline{a_n})$ describes what can be observed when running (under the constraints given by I) the observer functions $o \in \mathcal{O}$ over appropriate symbolic variables from $\overline{a_n}$. Each delivered equation is formed by equating each observer call to the (symbolic) value that the call returns. We require o to compute the same symbolic values at the end of all its symbolic execution branches in order to distill a (partial) observational abstraction or explanation for a given configuration in terms of o.

The consequent q of the axiom is the conjunction of $ret = v$, which specifies the return value v of the method m as recorded in the k cell of F, and the equations given by $explain(F, \overline{a_n})$, which in terms of the observers characterizes the final pattern F of the given branch. Note that the return value v could be either uninit or an initialized typed value that represents the return value for m under the conditions given by ϕ.

It is important to note that, in the axioms, the different equations in the antecedent (resp. consequent) of every implication formula are assumed to be run independently of each other under the same initial configuration. This is achieved by the *explain* algorithm by using the same initial state when considering the different observer functions to explain I and F. This avoids making any assumptions about function purity or side-effects. Depending on the boolean value of the abstract subsumption flag $\sharp$ in F (line 9), the synthesized axiom ax is directly added to the postcondition Q (when $\sharp$ is *false*) or to the conjunction $Q^\sharp$ (when $\sharp$ is true) that collects all candidate axioms extracted from branches that contain at least one node that was folded by abstract subsumption. Note that, due to the under-approximation introduced by abstract subsumption [3], there may be some behaviors (real trace fragments) beyond the abstract folded states that are not captured by the deployed symbolic abstract traces. Therefore, axioms in $Q^\sharp$ could have spurious instances and must be double-checked. We apply a post-processing refinement $refine(Q, Q^\sharp)$ which tries to build specialized (demonstrably correct) instances of the axioms in $Q^\sharp$ that can be added to Q, while getting rid of any $Q^\sharp$ axioms that remain overly general (i.e., that can have both true and false instances). A further subsumption checking over the resulting set of axioms is included in the refinement post-processing that purges the augmented Q from less general axioms.

When Algorithm 1 terminates, the generated contract is $< P, Q, \mathcal{L} >$ where the method precondition P is the disjunction of all axiom premises, the method postcondition is given by $refine(Q, Q^\sharp)$, and $\mathcal{L}$ records all program locations that are (potentially) modifiable by m. Note that we do not need to specialize the disjunction P according to the final refined postcondition Q because correctness of the contract is ensured by the specialized axiom guards of Q.

We note that lazy initialization is not applied when symbolically executing the observer functions. This is because we want to start from an initial configuration whose dynamic memory satisfies (or is given by) φ, and if any uninitialized addresses are expanded by lazy initialization, such an initial configuration (and thus the target of the observation) would be altered. This implies that some final patterns in the symbolic execution trees for the given observer may contain uninit return values, meaning that we know nothing regarding the dynamic memory from that point on. When this occurs

(for all branches), the *explain* algorithm generates a conjunct where the observer call is equated to a fresh symbolic value.

Let us compute a specification for the `insert` modifier function of Example 1 by applying Algorithm 1.

Example 5. We first compute $\mathrm{SE}^{\sharp}(\texttt{insert(\&s, ?x)\{true\}})$ with `&s` being a symbolic address with initial value `uninit` and with `?x` being a symbolic integer value. Since there are no constraints in the initial symbolic configuration, the execution covers all possible initial concrete configurations. Then, the abstract symbolic execution computes ten final configurations. The following one represents the final state for the path where the `while` loop stops due to abstract subsumption between the states associated to two consecutive iterations (Nodes S_{18} and S_{27} of Example 4):

$$\langle \mathrm{tv}(int, 1) \rangle_{\mathsf{k}}$$
$$\Big\langle \begin{array}{c} \texttt{s} \mapsto \texttt{\&s}, \texttt{x} \mapsto \texttt{?x}, \texttt{new_node} \mapsto \texttt{\&new_node}, \\ \texttt{end_node} \mapsto \texttt{\&s.elems.next.next}, \texttt{n} \mapsto \texttt{\&s.elems.next.next.next} \end{array} \Big\rangle_{\mathsf{env}}$$
$$\Big\langle \begin{array}{c} \texttt{\&s} \mapsto (\texttt{capacity} \mapsto \texttt{?s.capacity}, \texttt{size} \mapsto \texttt{?s.size} + 1, \texttt{elems} \mapsto \texttt{\&new_node}), \\ \texttt{\&new_node} \mapsto (\texttt{value} \mapsto \texttt{?x}, \texttt{next} \mapsto \texttt{\&s.elems}), \texttt{\&x} \mapsto \texttt{?x} \\ \texttt{\&s.elems} \mapsto (\texttt{value} \mapsto \texttt{?v}_0, \texttt{next} \mapsto \texttt{\&s.elems.next}), \\ \cdots \\ \texttt{\&s.elems.next.next.next} \mapsto (\texttt{value} \mapsto \texttt{?v}_3, \texttt{next} \mapsto \texttt{uninit}) \end{array} \Big\rangle_{\mathsf{heap}}$$
$$\Big\langle \begin{array}{c} \texttt{\&s} \mapsto (\texttt{capacity} \mapsto \texttt{?s.capacity}, \texttt{size} \mapsto \texttt{?s.size}, \texttt{elems} \mapsto \texttt{\&s.elems}), \\ \texttt{\&s.elems} \mapsto (\texttt{value} \mapsto \texttt{?v}_0, \texttt{next} \mapsto \texttt{\&s.elems.next}), \\ \cdots \\ \texttt{\&s.elems.next.next.next} \mapsto (\texttt{value} \mapsto \texttt{?v}_3, \texttt{next} \mapsto \texttt{uninit}) \end{array} \Big\rangle_{\mathsf{init-heap}}$$
$$\langle \texttt{?s.size} < \texttt{?s.capacity} \wedge \texttt{?v}_0 \neq \texttt{?x} \wedge \texttt{?v}_1 \neq \texttt{?x} \wedge \texttt{?v}_2 \neq \texttt{?x} \rangle_{\mathsf{path-condition}}$$

Roughly speaking, the execution of this path corresponds to the case when the element `x` (with symbolic value `?x`) is effectively inserted in a non-empty list that contains three elements. The return value (k cell) of the call `insert(&s, ?x)` is the integer 1 (standing for success); the symbolic (initial) value `?s.size` of the field `size` of `s` is increased by 1 and now the field `elems` of `s` points to an object `&new_node` with `value` `?x` as the first node of the set. For the sake of simplicity, we omit any cell components that are irrelevant for comprehension.

As a side effect of applying abstract subsumption to stop the `while` loop, the node pointed by the field `next` of the last object node is not null but `uninit`. This implies a loss of precision: the symbolic heap is matched by any concrete heap whose first node contains the value `?x` and is followed by 3 or more nodes.

The algorithm computes the explanation for the corresponding initial and final state of each of those ten configurations. Let us illustrate one of the cases.

Example 6. In order to explain the final pattern F of Example 5, the function *explain* considers the universe of observer calls, which include the call `contains(&s,?x)`. The symbolic execution of `contains(&s,?x)` under the constraints given by the heap and path-condition cells of F results in a single-branch tree with return value 1; hence, the equation `contains(s,x)=1` is added as part of the equational explanation of F.

Example 7. In order to explain the corresponding initial pattern I, we symbolically execute the observer `contains(&s,?x)` under the constraints given by I (i.e., the init-heap and path-condition cells of F); and since no element with value `?x` is found in the set `&s`, the list is traversed until the `uninit` node is reached. Hence, the equation `contains(s,x)=`$_v$ is generated, with $_v$ being a symbolic value that stands for any possible value that the function may return (either 0 or 1 in this example).

Finally, let us illustrate how the refinement process $\mathit{refine}(Q, Q^\sharp)$ for method m works. Roughly speaking, for each candidate axiom $p \Rightarrow q$ in $Q^\sharp$, we first randomly generate test cases (initial configurations) that satisfy the axiom antecedent p, then we run the modifier method m on the initial configurations, and finally we check whether or not the consequent q is satisfied after the method execution. Refuted candidate axioms are not automatically removed: a counterexample-guided, specialization post-process defined in [1] is attempted first. It uses the concrete values refuting the axiom (or more precisely the deployed symbolic execution branches resulting from fixing those values on the initial configuration), as counterexample behaviors to be excluded from the symbolic execution tree. Then, by iteratively repeating the inference process on the reduced tree, new axioms that are either eventually correct (and then added to Q) or can be further specialized are distilled. Note that this process is guaranteed to terminate since the size of the tree is reduced at each iteration.

Example 8. After the for `loop` of Algorithm 1, one axiom for each of the (10) final patterns is synthesized. After removing duplicates, 7 axioms are kept (see Figure 4), together with one candidate axiom (labelled as `C1`) that derives from the final configuration discussed in Example 5.

$$\texttt{A1}\ \begin{pmatrix}\texttt{isnull(s)} = 1\ \wedge \texttt{isempty(s)} = 0\ \wedge\\ \texttt{isfull(s)} = 0\ \wedge \texttt{contains(s,x)} = 0\ \wedge\\ \texttt{length(s)} = 0\end{pmatrix} \Rightarrow \begin{pmatrix}\texttt{isnull(s')} = 1\ \wedge \texttt{isempty(s')} = 0\ \wedge\\ \texttt{isfull(s')} = 0\ \wedge \texttt{contains(s',x)} = 0\ \wedge\\ \texttt{length(s')} = 0\ \wedge \texttt{ret} = 0\end{pmatrix}$$

$$\texttt{A2}\ \begin{pmatrix}\texttt{isnull(s)} = 0\ \wedge \texttt{isempty(s)} = \texttt{_}i1\ \wedge\\ \texttt{isfull(s)} = 1\ \wedge \texttt{contains(s,x)} = \texttt{_}i2\ \wedge\\ \texttt{length(s)} = \texttt{_}i3\end{pmatrix} \Rightarrow \begin{pmatrix}\texttt{isnull(s')} = 0\ \wedge \texttt{isempty(s')} = \texttt{_}i1\ \wedge\\ \texttt{isfull(s')} = 1\ \wedge \texttt{contains(s',x)} = \texttt{_}i2\ \wedge\\ \texttt{length(s')} = \texttt{_}i3\ \wedge \texttt{ret} = 0\end{pmatrix}$$

$$\texttt{A3}\ \begin{pmatrix}\texttt{isnull(s)} = 0\ \wedge \texttt{isempty(s)} = 0\ \wedge\\ \texttt{isfull(s)} = 0\ \wedge \texttt{contains(s,x)} = 1\ \wedge\\ \texttt{length(s)} = \texttt{_}i1\end{pmatrix} \Rightarrow \begin{pmatrix}\texttt{isnull(s')} = 0\ \wedge \texttt{isempty(s')} = 0\ \wedge\\ \texttt{isfull(s')} = 0\ \wedge \texttt{contains(s',x)} = 1\ \wedge\\ \texttt{length(s')} = \texttt{_}i1\ \wedge \texttt{ret} = 0\end{pmatrix}$$

$$\texttt{A4}\ \begin{pmatrix}\texttt{isnull(s)} = 0\ \wedge \texttt{isempty(s)} = 1\ \wedge\\ \texttt{isfull(s)} = 0\ \wedge \texttt{contains(s,x)} = 0\ \wedge\\ \texttt{length(s)} = 0\end{pmatrix} \Rightarrow \begin{pmatrix}\texttt{isnull(s')} = 0\ \wedge \texttt{isempty(s')} = 0\ \wedge\\ \texttt{contains(s',x)} = 1\ \wedge \texttt{length(s')} = 1\ \wedge\\ \texttt{ret} = 1\end{pmatrix}$$

$$\texttt{A5}\ \begin{pmatrix}\texttt{isnull(s)} = 0\ \wedge \texttt{isempty(s)} = 0\ \wedge\\ \texttt{isfull(s)} = 0\ \wedge \texttt{contains(s,x)} = 0\ \wedge\\ \texttt{length(s)} = 1\end{pmatrix} \Rightarrow \begin{pmatrix}\texttt{isnull(s')} = 0\ \wedge \texttt{isempty(s')} = 0\ \wedge\\ \texttt{contains(s',x)} = 1\ \wedge \texttt{length(s')} = 2\ \wedge\\ \texttt{ret} = 1\end{pmatrix}$$

$$\texttt{A6}\ \begin{pmatrix}\texttt{isnull(s)} = 0\ \wedge \texttt{isempty(s)} = 0\ \wedge\\ \texttt{isfull(s)} = 0\ \wedge \texttt{contains(s,x)} = 0\ \wedge\\ \texttt{length(s)} = 2\end{pmatrix} \Rightarrow \begin{pmatrix}\texttt{isnull(s')} = 0\ \wedge \texttt{isempty(s')} = 0\ \wedge\\ \texttt{contains(s',x)} = 1\ \wedge \texttt{length(s')} = 3\ \wedge\\ \texttt{ret} = 1\end{pmatrix}$$

$$\texttt{A7}\ \begin{pmatrix}\texttt{isnull(s)} = 0\ \wedge \texttt{isempty(s)} = 0\ \wedge\\ \texttt{isfull(s)} = 0\ \wedge \texttt{contains(s,x)} = 0\ \wedge\\ \texttt{length(s)} = 3\end{pmatrix} \Rightarrow \begin{pmatrix}\texttt{isnull(s')} = 0\ \wedge \texttt{isempty(s')} = 0\ \wedge\\ \texttt{contains(s',x)} = 1\ \wedge \texttt{length(s')} = 4\ \wedge\\ \texttt{ret} = 1\end{pmatrix}$$

$$\texttt{C1}\ \begin{pmatrix}\texttt{isnull(s)} = 0\ \wedge \texttt{isempty(s)} = 0\ \wedge\\ \texttt{isfull(s)} = 0\ \wedge \texttt{contains(s,x)} = \texttt{_}i1\ \wedge\\ \texttt{length(s)} = \texttt{_}i2\end{pmatrix} \Rightarrow \begin{pmatrix}\texttt{isnull(s')} = 0\ \wedge \texttt{isempty(s')} = 0\ \wedge\\ \texttt{contains(s',x)} = 1\ \wedge \texttt{length(s')} = \texttt{_}i2 + 1\ \wedge\\ \texttt{ret} = 1\end{pmatrix}$$

Figure 4. Set of axioms and candidates for Example 5.

The refinement process is then triggered over `C1` to check if it can first be falsified and then refined. Given the binary domain 0/1 of the `contains(s,x)` function, the axiom is straightforwardly falsified (e.g., by the test case where `s` is a non-full set containing a single element with value 5, and `x` is 5). The final state does not satisfy the postcondition of axiom `C1` because, since the set `s` already contained the desired element, the modifier `insert(s,x)` does not return 1 and the length does not increase after the execution.

Since the axiom has been falsified (with `contains(s,x)=1`), now the refinement process is run with $\texttt{_}i1 \mapsto 0$ and the last (specialized and correct) axiom is obtained:

$$\texttt{A8}\left(\begin{array}{l}\texttt{isnull(s)} = 0\ \wedge \texttt{isempty(s)} = 0\ \wedge\\ \texttt{isfull(s)} = 0\ \wedge \texttt{contains(s,x)} = 0\ \wedge\\ \texttt{length(s)} = _i1\end{array}\right) \Rightarrow \left(\begin{array}{l}\texttt{isnull(s}') = 0\ \wedge \texttt{isempty(s}') = 0\ \wedge\\ \texttt{contains(s}', \texttt{x}) = 1\ \wedge \texttt{length(s}') = _i1 + 1\ \wedge\\ \texttt{ret} = 1\end{array}\right)$$

Note that the new axiom subsumes the fifth, sixth, and seventh axioms of the previous specification; hence, they are removed. After the refinement, the contract postcondition returned for `insert(s,x)` contains five axioms, specifically the axioms `A1-4` and `A8`.

As for the last element of the contract, the set of assignable program locations $\mathcal{L}$ is obtained as the union of the location sets that are recorded in the $\langle\rangle_{\mathsf{locations}}$ cells of the final symbolic execution states, which is $\mathcal{L} = \{\texttt{s}, \texttt{end_node}, \texttt{n}, \texttt{new_node}, \texttt{new_node} \mapsto \texttt{value}, \texttt{new_node} \mapsto \texttt{next}, \texttt{s} \mapsto \texttt{elems}, \texttt{s} \mapsto \texttt{size}\}$.

6 Related work and Conclusions

The wide interest in formal specifications as helpers for other analysis, validation, and verification tools have resulted in numerous approaches for (semi-)automatically computing different kinds of specifications that can take the form of contracts, snippets, summaries, properties, process models, rules, graphs, automata, interfaces, or component abstractions.

Let us briefly discuss those strands of research that have influenced our work the most. A detailed description of the related literature can be found in [22,1,8]. Our axiomatic representation is inspired by [21], which relies on a model checker for symbolic execution and generates either Spec# specifications or parameterized unit tests. In contrast to [21], we take advantage of $\mathbb{K}$ symbolic capabilities to generate simpler and more accurate formulas that avoid reasoning with the global heap because the different pieces of the heap that are reachable from the function argument addresses are kept separate. Unlike our symbolic approach, QuickSpec [7], Daikon [11], and the algebraic specification discovery tool of Henkel and Diwan [13] detect program assertions by extensive testing. Whereas Daikon discovers invariants that hold at existing program points, QuickSpec discovers equations between arbitrary terms (laws) that are constructed by using an API. This is similar to the approach of Henkel and Diwan [13], which generalizes the results of running tests on Java class interfaces as an algebraic specification. By combining the concrete execution of actual test cases with a simultaneous symbolic execution of the same tests, DySy determines program properties that generalize the observed behaviors [9]. Starting from simple, partial contracts previously written by the programmer, rich post-conditions involving quantification are defined in [22] by using random testing. Other approaches to software specification discovery based on abstract interpretation are [20,5,8], while [23,12] use inductive matching learning.

Since our approach is generic and not tied to the $\mathbb{K}$ semantics specification of KernelC, we expect that the methodology developed in this work can be easily extended to other languages for which a $\mathbb{K}$ semantics is given. Moreover, the correctness of the delivered specifications can be automatically ensured by using the existing $\mathbb{K}$ formal tools.

We have developed a prototype implementation of the extended $\mathbb{K}$ symbolic machinery and contract inference algorithm described in the previous sections (available at `http://safe-tools.dsic.upv.es/kindspec2`), and we have used it to mechanize our runnning example. The abstraction component is not fully integrated within the system, yet it can be used by manually fixing the abstract domain for the program at hand.

Our preliminary results are promising since they show that general correct axioms can be inferred, leading to a more compact, clear, and complete specification. The contracts generated by our tool can be easily translated to richer (but also heavier) notations for behavioural interface C specifications such as ACSL or to the native syntax of some SMT solvers, which is planned as future work.

References

1. Alpuente, M., Feliú, M.A., Villanueva, A.: Automatic Inference of Specifications Using Matching Logic. Proc. of PEPM'13, 127–136. ACM (2013)
2. Alpuente, M., Pardo, D., Villanueva, A.: Automatic Inference of Specifications in the K Framework. EPTCS 200: 1–17 (2015)
3. Anand, S., Păsăreanu, C.S., Visser, W.: Symbolic execution with abstraction. STTT 11(1): 53–67 (2008)
4. Arusoaie, A., Lucanu, D., Rusu, V.: Symbolic execution based on language transformation. Computer Languages, Systems & Structures 44, Part A, 48–71 (2015)
5. Bacci, G., Comini, M., Feliú, M.A., Villanueva, A.: Automatic Synthesis of Specifications for First Order Curry Programs. Proc. PPDP'12, 25–34. ACM (2012)
6. Baudin, P., Filliâtre, J.C., Marché, C., Monate, B., Moy, Y., Prevosto, V.: ACSL: ANSI/ISO C Specification Language, version 1.4 (2009)
7. Claessen, K., Smallbone, N., Hughes, J.: QuickSpec: Guessing Formal Specifications Using Testing. Tests and Proofs 2010, Springer LNCS 6143: 6–21 (2010)
8. Cousot, P., Cousot, R., Fähndrich, M., Logozzo, F.: Automatic Inference of Necessary Preconditions. Proc. of VMCAI'13, Springer LNCS 7737: 128–148 (2013)
9. Csallner, C., Tillmann, N., Smaragdakis, Y.: DySy: Dynamic Symbolic Execution for Invariant Inference. Proc. of ICSE'08, 281–290. ACM (2008)
10. Ellison, C., Roşu, G.: An Executable Formal Semantics of C with Applications. Proc. of POPL'12, 533–544. ACM (2012)
11. Ernst, M.D., Perkins, J.H., Guo, P.J., McCamant, S., Pacheco, C., Tschantz, M.S., Xiao, C.: The Daikon system for dynamic detection of likely invariants. SCP 69(1–3): 35–45 (2007)
12. Giannakopoulou, D., Păsăreanu, C.S.: Interface Generation and Compositional Verification in JavaPathfinder. Proc. of FASE'09, Springer LNCS 5503: 94–108 (2009)
13. Henkel, J., Diwan, A.: Discovering Algebraic Specifications from Java Classes. Proc. of ECOOP'03. Springer LNCS 2743: 431–456 (2003)
14. Khurshid, S., Păsăreanu, C.S., Visser, W.: Generalized Symbolic Execution for Model Checking and Testing. Proc. of TACAS'03, Springer LNCS 2619: 553–568 (2003)
15. King, J.C.: Symbolic Execution and Program Testing. Comm. ACM 19(7): 385–394 (1976)
16. Liskov, B., Guttag, J.: Abstraction and Specification in Program Development. MIT Press, Cambridge, MA, USA (1986)
17. Meyer, B.: Applying 'design by contract'. Computer 25(10): 40–51 (1992)
18. Moura, L.d., Bjørner, N.: Z3: An Efficient SMT Solver. Proc. of TACAS'08, Springer LNCS 4963: 337–340 (2008)
19. Roşu, G., Şerbănuţă, T.F.: An overview of the K semantic framework. JLAP 79(6): 397–434 (2010)
20. Taghdiri, M., Jackson, D.: Inferring specifications to detect errors in code. Automated Software Eng. 14(1): 87–121 (2006)
21. Tillmann, N., Chen, F., Schulte, W.: Discovering Likely Method Specifications. Proc. of ICFEM'06, Springer LNCS 4260: 717–736 (2006)
22. Wei, Y., Furia, C.A., Kazmin, N., Meyer, B.: Inferring Better Contracts. Proc. of the ICSE'11, 191–200. ACM (2011)
23. Whaley, J., Martin, M.C., Lam, M.S.: Automatic Extraction of Object-oriented Component Interfaces. Proc. of ISSTA'02, 218–228. ACM (2002)

A Full example code

```
#include <stdlib.h>

struct lnode{
 int value;
 struct lnode *next;
};

struct set {
 int capacity;
 int size;
 struct lnode *elems;
};

struct set* new(int capacity) {
 struct set *new_set;

 new_set = (struct set*) malloc(sizeof(struct set));
 if(new_set == NULL)
  return NULL; /* no memory left */

 new_set->capacity = capacity;
 new_set->size = 0;
 new_set->elems = NULL;
 return new_set;
}

int insert(struct set *s, int x) {
 struct lnode *new_node;
 struct lnode *end_node;
 struct lnode *n;

 if(s==NULL)
  return 0; /* NULL set */

 if(s->size >= s->capacity)
  return 0; /* no space left */

 end_node = s->elems;
 n = end_node;
 while(n != NULL) {
  if(n->value == x)
   return 0; /* element already in the set */
  end_node = n;
  n = n->next;
 }

 /* Creation of new node */
 new_node = (struct lnode*) malloc(sizeof(struct lnode));
 if(new_node == NULL)
  return 0; /* no memory left */
 new_node->value = x;
 new_node->next = s->elems;

 s->elems = new_node;
 s->size += 1;

 return 1; /* element added */
}

int isnull(struct set *s) {
 if(s==NULL)
  return 1;
 return 0;
}

int isempty(struct set *s) {
 if(s==NULL)
  return 0;
 if(s->elems==NULL)
  return 1; /* s is empty */
 return 0;
}

int isfull(struct set *s) {
 if(s==NULL)
  return 0;
 if(s->size >= s->capacity)
  return 1; /* s is full */
 return 0;
}

int contains(struct set *s, int x) {
 struct lnode *n;

 if(s==NULL)
  return 0; /* s is NULL */

 n = s->elems;
 while(n != NULL){
  if(n->value == x)
   return 1; /* element found */
  n = n->next;
 }

 return 0; /* element NOT found */
}

int length(struct set *s) {
 struct lnode *n;
 int count;

 if(s==NULL)
  return 0; /* s is NULL */

 count = 0;
 n = s->elems;
 while(n != NULL){
  count = count + 1;
  n = n->next;
 }

 return count;
}
```